

Approximation Algorithms for Secondary Spectrum Auctions*

Martin Hoefer

Thomas Kesselheim

Berthold Vöcking

Abstract

We study combinatorial auctions for the secondary spectrum market. In this market, short-term licenses shall be given to wireless nodes for communication in their local neighborhood. In contrast to the primary market, channels can be assigned to multiple bidders, provided that the corresponding devices are well separated such that the interference is sufficiently low. Interference conflicts are described in terms of a conflict graph in which the nodes represent the bidders and the edges represent conflicts such that the feasible allocations for a channel correspond to the independent sets in the conflict graph.

In this paper, we suggest a novel LP formulation for combinatorial auctions with conflict graph using a non-standard graph parameter, the so-called *inductive independence number*. Taking into account this parameter enables us to bypass the well-known lower bound of $\Omega(n^{1-\epsilon})$ on the approximability of independent set in general graphs with n nodes (bidders). We achieve significantly better approximation results by showing that interference constraints for wireless networks yield conflict graphs with bounded inductive independence number.

Our framework covers various established models of wireless communication, e.g., the protocol or the physical model. For the protocol model, we achieve an $O(\sqrt{k})$ -approximation, where k is the number of available channels. For the more realistic physical model, we achieve an $O(\sqrt{k} \log^2 n)$ approximation based on edge-weighted conflict graphs. Combining our approach with the LP-based framework of Lavi and Swamy, we obtain incentive compatible mechanisms for general bidders with arbitrary valuations on bundles of channels specified in terms of demand oracles.

1 Introduction

A major challenge of today's wireless networks and mobile communication is spectrum management, as devices use common frequency bands that are subject to interference between multiple transmitters in the same area. In fact, spectrum allocation has become one of the key problems that currently limits the growth and evolution of wireless networks. The reason is that, traditionally, frequencies were given away to large service providers in a static way by regulators for entire countries. Examples include FCC auctions in the US or the auctions for UMTS and LTE that took place in Europe. However, demands for services vary at different times and in different areas. Depending on time and place this causes frequency bands licensed for one application to become overloaded. On the other hand, different bands are idle at the same time. A promising solution to this problem is to use market approaches that result in a flexible and thus more efficient redistribution of access rights – thereby overcoming the artificial shortage of available spectrum. In this

*Department of Computer Science, RWTH Aachen University, Germany. {mhoefer,kesselheim,voecking}@cs.rwth-aachen.de. This work has been supported by DFG through UMIC Research Centre, RWTH Aachen University, and grant Ho 3831/3-1.

case, parts of the spectrum that are currently unused by so-called *primary users* for the originally intended purpose (such as TV or telecommunication) can be offered to so-called *secondary users*. Licenses for such secondary usage are valid only for a local area.

A sustainable approach (concisely termed “eBay in the Sky” in [33]) to automatically run such a secondary spectrum market is to auction licenses for secondary users on a regular basis. In this paper, we propose a general framework and efficient algorithms to implement such a secondary spectrum auction. In our model, there are n secondary users who can bid for bundles of the k wireless channels. Depending on the scenario a user can correspond to a base station that strives to cover a specific area or a pair of devices that want to exchange data (e.g., a base station and a mobile device). In order to account for channel aggregation capabilities of modern devices, users should be able to acquire multiple channels. We allow each user v to have an arbitrary valuation $b_{v,T}$ for each subset T of channels. This level of generality is necessary because of different needs, applications, and hardware abilities of the users, but also because of different locations, spectrum availability, and interference conditions. For instance, the presence of a primary user might allow access to a channel only for a subset of mobile devices located in selected areas. We assume no restrictions on the valuation functions, not even monotonicity.

In this paper, we devise approximation algorithms for spectrum allocation on the secondary market with the objective of maximizing social welfare. We focus on the underlying combinatorial problems and describe interference conflicts by an edge-weighted conflict graph. In unweighted graphs, the vertices represent the bidders and the edges represent conflicts such that the feasible allocations for a channel correspond to the independent sets in the conflict graph. For edge-weighted graphs, we extend the definition of independent set to weighted edges by requiring the sum of all incoming weights to be less than 1. We address the following problem.

Problem 1 (Combinatorial Auction with Conflict Graph). *Given a graph $G = (V, E)$, a natural number k , and a valuation function $b: V \times 2^{[k]} \rightarrow \mathbb{N}$, find a feasible allocation $S: V \rightarrow 2^{[k]}$ that maximizes the social welfare $b(S) := \sum_{v \in V} b_{v,S(v)}$.*

An allocation S is called feasible if for all channels $j \in [k]$, the set of vertices that are assigned to this channel, i. e. $\{v \in V \mid j \in S(v)\}$, is an independent set.

Observe that this problem generalizes combinatorial auctions (where the conflict graph is a clique) and maximum weight independent set (where $k = 1$). This formulation covers a large number of binary interference models (such as the protocol model). As we will see, edge weights allow to express even more realistic models like the physical model. Here, we can even take the effects of power control into account.

1.1 Our contribution

We devise the first approximation algorithms for the combinatorial auction problem with conflict graph. Our approach is based on a novel LP formulation for the independent set problem using a non-standard graph parameter.

Definition 1 (inductive independence number ρ). *For a graph $G = (V, E)$, the inductive independence number ρ is the smallest number such that there is an ordering π of the vertices satisfying: For all $v \in V$ and all independent sets $M \subseteq V$, we have $|M \cap \{u \in V \mid \{u, v\} \in E, \pi(u) < \pi(v)\}| \leq \rho$.*

In words, for every vertex $v \in V$, the size of an independent set in the *backward neighborhood* of v , i.e., the set of neighbors u of v with $\pi(u) < \pi(v)$, is at most ρ . Conflict graphs derived from

various simple models of wireless communication with binary conflicts like, e.g., the protocol model, distance-2 matchings, or disk graphs, have $\rho = O(1)$, see, e.g., [31]. The corresponding ordering π is efficiently computable in these cases. We exploit this property in our algorithms.

Our main results concern the so-called physical model which is common in the engineering community and was only recently subject to theoretical work. In binary models of wireless communication usually studied in theoretical computer science, we make the oversimplifying assumption that interference caused by a signal stops at some boundary around the sender, and receivers beyond this boundary are not disturbed by this signal. In contrast, the physical model takes into account realistic propagation effects and additivity of signals. Feasibility of simultaneous transmissions is modeled in terms of signal to interference plus noise ratio (SINR) constraints. We study two variants of this model, one in which signals are sent at given powers (e.g., uniform) and one where the powers are subject to optimization themselves. We show how to represent SINR constraints for both of these variants in terms of an edge-weighted conflict graph and introduce appropriate notions of “independent set” and “inductive independence number” for edge-weighted graphs. Note that the combinatorial auctions with edge-weighted conflict graphs can be defined in the same way as stated in Problem 1 given an appropriate definition of “independent set”.

At first, we prove that the inductive independence number ρ for edge-weighted graphs obtained from the physical model (in both variants) is bounded by $O(\log n)$ and the corresponding ordering is efficiently computable. This enables us to bypass the well-known lower bound of $\Omega(n^{1-\epsilon})$ on the approximability of independent set in general graphs. In particular, we present an LP relaxation capturing both interference constraints and valuations of users for subsets of channels. Similar to regular combinatorial auctions, the LP might require an exponential number of valuations $b_{v,T}$ to be written down explicitly. However, we show how to solve the LP using only oracle access to bidder valuations. Our LP based framework is able to handle edge-weighted conflict graphs resulting from the physical model. By rounding the LP optimum, our algorithm achieves an $O(\rho \cdot \sqrt{k} \log n)$ approximation guarantee. Combining this with the bound on ρ gives an $O(\sqrt{k} \log^2 n)$ -approximation of the social welfare for spectrum auctions in the physical model (in both variants).

For more simple binary models of wireless communication such as the protocol model, our approach yields an $O(\rho \cdot \sqrt{k})$ -approximation. Using the bounds on ρ mentioned above, this yields an $O(\sqrt{k})$ approximation guarantee. In this case, we also provide some complementing hardness results. In general, it is hard to approximate the combinatorial auction problem with conflict graphs to a factor of $O(\rho^{1-\epsilon})$ and to a factor of $O(k^{\frac{1}{2}-\epsilon})$ for any constant $\epsilon > 0$. While for some specific models better approximations exist, in general the bounds provided by our algorithms for binary models cannot be improved in terms of a single parameter ρ or k . In addition, we provide stronger lower bounds for the case of asymmetric channels, in which the conflict graph can be different for each channel. In this case, our algorithm guarantees a factor of $O(\rho \cdot k)$, which is best possible in general.

Our approach can be used to derive incentive compatible mechanisms using the LP-based framework of Lavi and Swamy [27] for general bidders with demand oracles. In fact, we slightly extend this framework by starting with an infeasible rather than feasible ILP formulation. The approximation algorithm computes a linear combination of feasible solutions approximating the optimal solution of the corresponding LP and then chooses one of these solutions at random. The obtained mechanism is truthful in expectation.

Outline. For technical reasons, we present our results in a different order than stated above. We first introduce the basic approach in the context of unweighted conflict graphs in Section 2. The

extensions to edge-weighted graphs including formal definitions of independent sets and inductive independence number are given in Section 3. The aforementioned wireless models (especially the variants of the physical model) are formally introduced in Section 4, where we also show the bounds on the inductive independence number. The application of the framework by Lavi and Swamy is discussed in Section 5. Finally, the results on asymmetric channels are presented in Section 6.

1.2 Related Work

The idea of establishing secondary spectrum markets has attracted much attention among researchers in applied networking and engineering communities [33, 14, 5, 22]. There are many different fundamental regulatory questions that need to be addressed when implementing such a market. For example it has to be clarified who runs the market and who is allowed to sell and buy spectrum there. Possible actors could be network providers, brokers, regulators and end-users. In addition, it has to be guaranteed that existing services are not harmed. In most of the literature on spectrum markets the technological aspects dominate. Many results in this area are only of qualitative nature, only a few examples (such as [33, 34]) do explicitly consider truthfulness. We believe that our combinatorial models based on (edge-weighted) conflict graphs taking into account the bounded inductive independence number allows us to neglect technological aspects and to focus on the underlying combinatorial and algorithmic questions. To the best of our knowledge there is no previous work on auctions using the general framework of conflict graphs, or, in general, non-trivial provable worst-case guarantees on the efficiency of the allocation.

In contrast, combinatorial auctions have been a prominent research area in algorithmic game theory over the last decade. A variety of works treats auctions with special valuation functions, such as submodular valuations or ones expressible by specific bidding languages. For an introduction see, e.g., [28, Chapters 11 and 12] or [7]. In addition, designing (non-truthful) approximation algorithms for the allocation problems has found interest, most notably for submodular valuations (e.g., [30, 12]). More relevant to our work, however, are results that deal with truthful mechanisms for general valuations. Most notably, Lavi and Swamy [27] and Dobzinski et al. [8] derive mechanisms using only demand oracles that achieve an \sqrt{k} -approximation with truthfulness in expectation and universal truthfulness, respectively. A deterministic truthful $(k/\sqrt{\log k})$ -approximation is obtained by Holzman et al. [21].

Over the last decades, there has been much research on finding maximum independent sets in the context of interference models for wireless networks. One of the simplest models in this area are disk graphs, which are mostly analyzed using geometric arguments. See [13, 16] for a summary on the results and typical techniques. Recently and independently from our work, Christodoulou et al. [6] study combinatorial auctions for geometric objects. Similar to our approach, they present an LP formulation based on a property in terms of an ordering, the fatness of geometric objects.

Akcoglu et al. [1] and Ye and Borodin [32] also use the inductive independence number to approximate independent sets within a factor of ρ with a motivation stemming from chordal graphs. However, they do not consider multiple channels or wireless communication. As the algorithm is not monotone, it is also not immediately applicable for truthful auctions.

Algorithmic aspects of the physical model have become popular in theoretical research recently, particularly the problem of scheduling, i.e., partitioning a given set of requests in a small number of classes such that all requests are successful. New challenges arise since graph-theoretic coloring methods cannot be directly applied. For example, there have been a number of results on how to choose powers for short schedule lengths [11, 10, 18]. A popular method is fixing powers according

to some distance-based scheme. For uniform power assignments, a constant-factor approximation algorithm for the problem of finding an independent set (i.e., a maximum set that may share a single channel) is presented in [15]. An online version of the problem has been studied in [9] presenting tight bounds depending on the difference in lengths of the requests. Most recently, a constant-factor approximation algorithm for arbitrary power schemes has been obtained by Kesselheim [24].

2 Unweighted Conflict Graphs

2.1 Our LP relaxation

One can get a very intuitive LP formulation for the Weighted Independent Set problem by leaving out the integer constraints from the Integer Linear Programm formulation.

$$\begin{aligned}
& \text{Max.} && \sum_{v \in V} b_v x_v \\
& \text{s. t.} && x_u + x_v \leq 1 && \text{for all } \{u, v\} \in E \\
& && 0 \leq x_v \leq 1 && \text{for all } v \in V
\end{aligned}$$

This LP can be used to approximate Independent Set within a factor of $(\bar{d} + 1)/2$ [20, 23] where \bar{d} is the average vertex degree. However, even for the case of a clique the integrality gap is $n/2$.

In contrast to this edge-based LP formulation, we here present a different LP based on the *inductive independence number* ρ (recall Definition 1). As we will see later, in typical conflict graphs the inductive independence number is constant and the corresponding ordering π can be efficiently calculated. Here we use $\Gamma_\pi(v) = \{u \in V \mid \{u, v\} \in E, \pi(u) < \pi(v)\}$ to denote the backward neighborhood of v . This allows to use the following LP relaxation that has one constraint for each combination of a vertex and a channel and another one for each vertex.

$$\text{Max.} \quad \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x_{v,T} \tag{1a}$$

$$\text{s. t.} \quad \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T \subseteq [k] \\ j \in T}} x_{u,T} \leq \rho \quad \text{for all } v \in V, j \in [k] \tag{1b}$$

$$\sum_{T \subseteq [k]} x_{v,T} \leq 1 \quad \text{for all } v \in V \tag{1c}$$

$$x_{v,T} \geq 0 \quad \text{for all } v \in V, T \subseteq [k] \tag{1d}$$

This LP works as follows. For each vertex v and each possible set $T \subseteq [k]$ of channels assigned to this vertex, there is one variable $x_{v,T}$. Due to the bounded inductive independence number all feasible allocations correspond to solutions of the LP. However, not all integer solutions of the LP necessarily correspond to feasible channel allocations. Nevertheless, we will show how to compute a feasible allocation from each solution.

Lemma 1. *Let S be a feasible allocation and x be defined by $x_{v,T} = 1$ if $S(v) = T$ and 0 otherwise, then x is a feasible LP solution.*

Proof. Conditions (1c) and (1d) are obviously satisfied. Let us now consider Condition (1b) for some fixed $v \in V$, $j \in [k]$. Set $M := \{u \in V \mid \pi(u) < \pi(v), j \in S(u)\}$. Since M is an independent set, by definition of the inductive independence number, we have $|M \cap \Gamma_\pi(v)| \leq \rho$.

On the other hand, we have

$$\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T \subseteq [k] \\ j \in T}} x_{u,T} = |M \cap \Gamma_\pi(v)| \leq \rho .$$

So x is a feasible LP solution. \square

As all coefficients are non-negative, this LP has a packing structure. In particular, we can observe the following decomposition property.

Observation 2. *Let x be a feasible solution to the LP, and $x^{(1)}$ be a vector such that $0 \leq x_{v,T}^{(1)} \leq x_{v,T}$ for all $v \in V$, $T \subseteq [k]$. Then $x^{(1)}$ and $x^{(2)} := x - x^{(1)}$ are feasible LP solutions as well.*

If there are only $\mathcal{O}(\log n)$ valuations $b_{v,T}$ non-zero, this LP is solvable in polynomial time. In general, the elementary representation of the $b_{v,T}$ values is exponential in k . We can still solve the LP optimally if bidders can be represented by demand oracles.

2.2 Demand Oracles

If there is an arbitrary number of channels, we must define an appropriate way to query the valuation functions of the requests, as an elementary description becomes prohibitively large. A standard way to deal with this issue in the auction literature is the representation by so-called *demand oracles*. To query the demand oracle of bidder v , we assign each channel i a price p_i . Then the oracle delivers his “demand” $S = \arg \max_{T \subseteq [k]} b_{v,T} - \sum_{i \in T} p_i$, i. e., a bundle that maximizes the utility of v given that he pays the sum of prices of channels in the bundle. In ordinary combinatorial auctions such demand oracles can be used to separate the dual of the underlying LP. We here show that such demand oracles can also be used for the solution of our LP (1). Consider the dual given by

$$\text{Min. } \sum_{v \in V} \sum_{j \in [k]} \rho y_{v,j} + \sum_{v \in V} z_v \tag{2a}$$

$$\text{s. t. } \sum_{\substack{u \in V \\ v \in \Gamma_\pi(u)}} \sum_{j \in T} y_{u,j} + z_v \geq b_{v,T} \quad \text{for all } v \in V, T \subseteq [k] \tag{2b}$$

$$y_{v,j} \geq 0 \quad \text{for all } v \in V, j \in T \tag{2c}$$

In contrast to ordinary combinatorial auctions, we cannot use the solution (y, z) directly as the channel prices. Instead, we choose *bidder-specific* channel prices by

$$p_{v,j} = \sum_{\substack{u \in V \\ v \in \Gamma_\pi(u)}} y_{u,j} .$$

Using this idea we see that the constraints of the dual are indeed equivalent to upper bounds on the utility with bidder-specific channel prices. By obtaining the demand bundle with highest utility

for each player, we find a violated constraint or verify that none exists. This allows to separate the dual LP and to solve it efficiently using the ellipsoid method. This way, we get an equivalent primal LP with only polynomially constraints. The corresponding primal solution has only polynomially many variables with $x_{v,T}^* > 0$.

2.3 Rounding LP Solutions

Having described the LP relaxation, we now analyze Algorithm 1 computing feasible allocations from LP solutions as follows. First, it decomposes the given LP solution to two solutions $x^{(1)}$ and $x^{(2)}$ (line 1). In $x^{(1)}$ all fractional variables $x_{v,T}$ for sets T with $|T| \geq \sqrt{k}$ are set to zero. To get $x^{(2)}$ the exact opposite is performed. From each one, a feasible allocating is computed and the better one is selected at the end. This means, the algorithm either allocates only sets of size at most \sqrt{k} or only of size at least \sqrt{k} . The actual computation of the allocation works the same way for both LP solutions. It consists of two major parts: a rounding stage and a conflict-resolution stage. In the rounding stage (lines 3–4), a tentative allocation is generated as follows. For each vertex v the set of allocated channels $S^{(l)}(v)$ is determined independently at random. Each set $T \neq \emptyset$ is taken with probability $x_{v,T}^{(l)}/2\sqrt{k}\rho$ and with the remaining probability the empty set is allocated.

Conflicts can occur when two adjacent vertices share the same channel. In this case, the conflict is resolved (lines 5–8) by allocating the channel to the vertex with smaller index in the π ordering. The other vertex is removed from the solution by being allocated the empty set.

Algorithm 1: LP rounding algorithm for the combinatorial auction problem with unweighted conflict graphs

```

1  decompose  $x$  into two solutions  $x^{(1)}$  and  $x^{(2)}$  by  $x_{v,T}^{(1)} = x_{v,T}$  if  $|T| \leq \sqrt{k}$  and  $x_{v,T}^{(1)} = 0$ 
   otherwise.  $x^{(2)} = x - x^{(1)}$  ;
2  for  $l \in \{1, 2\}$  do
3      for  $v \in V$  do                                     /* Rounding Stage */
4           $\lfloor$  with probability  $\frac{x_{v,T}^{(l)}}{2\sqrt{k}\rho}$  set  $S^{(l)}(v) := T$ ;
5      for  $v \in V$  do                                     /* Conflict-Resolution Stage */
6          for  $u \in V$  with  $\pi(u) < \pi(v)$  and  $\{u, v\} \in E$  do
7              if  $S^{(l)}(u) \cap S^{(l)}(v) \neq \emptyset$  then
8                   $\lfloor$   $S^{(l)}(v) := \emptyset$ ;
9  return the better one of the solutions  $S^{(1)}$  and  $S^{(2)}$ ;

```

Theorem 3. *For any feasible LP solution x^* with value b^* , Algorithm 1 calculates a feasible allocation S of value at least $b^*/8\sqrt{k}\rho$ in expectation.*

Proof. The allocations $S^{(1)}$ and $S^{(2)}$ are obviously feasible allocations because if $\{u, v\} \in E$, then $S^{(1)}(u) \cap S^{(1)}(v) = \emptyset$ and $S^{(2)}(u) \cap S^{(2)}(v) = \emptyset$. Therefore, the output is also a feasible allocation.

Let us now bound the expected values of solutions $S^{(1)}$ and $S^{(2)}$. Let $l \in \{1, 2\}$ be fixed. Let $X_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the rounding stage. Clearly,

we have

$$\mathbf{E}[X_{v,T}] = \frac{x_{v,T}^{(l)}}{2\sqrt{k}\rho} . \quad (3)$$

Let $X'_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the conflict-resolution stage. We consider the event that $X'_{v,T} = 0$, given that $X_{v,T} = 1$, i.e. that v is removed in the conflict-resolution stage after having survived the rounding stage.

Lemma 4. *The probability of being removed in the conflict-resolution stage after having survived the rounding stage is at most $1/2$.*

Proof. The event can only occur if $X_{u,T'} = 1$ for some $u \in V$ with $\pi(u) < \pi(v)$, $\{u, v\} \in E$, and $T \cap T' \neq \emptyset$. In terms of the random variables $X_{u,T}$ this is

$$\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \geq 1 .$$

Using this notation we can bound the probability of the event by using the Markov inequality

$$\begin{aligned} \Pr[X'_{v,T} = 0 \mid X_{v,T} = 1] &\leq \Pr \left[\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \geq 1 \right] \\ &\leq \mathbf{E} \left[\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \right] . \end{aligned}$$

We will now show separately that this expectation is at most $1/2$ for each of the two possible values of l ($l = 1$ or $l = 2$).

Case 1 ($l = 1$): We have:

$$\mathbf{E} \left[\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \right] \leq \mathbf{E} \left[\sum_{j \in T} \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} X_{u,T'} \right] .$$

Due to linearity of expectation this is equal to

$$\sum_{j \in T} \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} \mathbf{E}[X_{u,T'}] .$$

Using Equation (3) and the fact that $x^{(1)}$ is an LP solution, this is

$$\sum_{j \in T} \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} \frac{x_{u,T'}^{(1)}}{2\sqrt{k}\rho} \leq \sum_{j \in T} \frac{1}{2\sqrt{k}} .$$

Recall that we only have to deal with sets T for which $|T| \leq \sqrt{k}$ in this case. Hence, the expectation is at most $1/2$, and so is the probability that v is removed in the conflict-resolution stage.

Case 2 ($l = 2$): In this case, we have $X_{u,T'} > 0$ only for sets T' with $|T'| \geq \sqrt{k}$. This yields for all $u \in V$

$$\sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \leq \sum_{\substack{T' \subseteq [k] \\ T' \neq \emptyset}} X_{u,T'} = \sum_{\substack{T' \subseteq [k] \\ T' \neq \emptyset}} \sum_{j \in T'} \frac{X_{u,T'}}{|T'|} = \sum_{j \in [k]} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} \frac{X_{u,T'}}{|T'|} \leq \frac{1}{\sqrt{k}} \sum_{j \in [k]} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} X_{u,T'} .$$

So, we get

$$\mathbf{E} \left[\sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} X_{u,T'} \right] \leq \mathbf{E} \left[\frac{1}{\sqrt{k}} \sum_{j \in [k]} \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} X_{u,T'} \right] .$$

Again, we use linearity of expectation, Equation (3) and the fact that $x^{(2)}$ is an LP solution. This gives us

$$\frac{1}{\sqrt{k}} \sum_{j \in [k]} \sum_{u \in \Gamma_\pi(v)} \sum_{\substack{T' \subseteq [k] \\ j \in T'}} \frac{x_{u,T'}^{(2)}}{2\sqrt{k}\rho} \leq \frac{1}{\sqrt{k}} \sum_{j \in [k]} \frac{1}{2\sqrt{k}} \leq \frac{1}{2} .$$

This bounds the probability for the second case.

In both cases we have $\mathbf{Pr} [X'_{v,T} = 0 \mid X_{v,T} = 1] \leq 1/2$. □

Using Lemma 4 and Equation (3) we get for all $v \in V$ and $T \subseteq [k]$

$$\mathbf{E} [X'_{v,T}] \geq \frac{x_{v,T}^{(l)}}{4\sqrt{k}\rho} .$$

This yields that both calculated solutions $S^{(l)}$ for $l \in \{1, 2\}$ have expected value

$$\mathbf{E} [b(S^{(l)})] = \mathbf{E} \left[\sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} \cdot X'_{v,T} \right] = \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} \cdot \mathbf{E} [X'_{v,T}] \geq \frac{1}{4\sqrt{k}\rho} \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x_{v,T}^{(l)} .$$

So, the better ones of the two solutions has expected value

$$\begin{aligned} \mathbf{E} [\max\{b(S^{(1)}), b(S^{(2)})\}] &\geq \frac{1}{2} \left(\mathbf{E} [b(S^{(1)})] + \mathbf{E} [b(S^{(2)})] \right) \\ &\geq \frac{1}{8\sqrt{k}\rho} \sum_{v \in V} \sum_{S \subseteq [k]} b_{v,S} (x_{v,S}^{(1)} + x_{v,S}^{(2)}) \\ &= \frac{1}{8\sqrt{k}\rho} \sum_{v \in V} \sum_{S \subseteq [k]} b_{v,S} x_{v,S} = \frac{b^*}{8\sqrt{k}\rho} . \end{aligned}$$

□

2.4 Hardness Results

In this section we provide matching lower bounds for the approximation ratios of our algorithms. This shows that the above results cannot be vitally improved without further restricting the model. Our results are based on the hardness of approximating independent set in bounded-degree graphs [29] or general graphs [19]. A first result is that the $\mathcal{O}(\rho)$ algorithm for the case $k = 1$ is almost optimal.

Theorem 5. *For $k = 1$ and for each $\rho = \mathcal{O}(\log n)$ there is no $\rho/2^{\mathcal{O}(\sqrt{\log \rho})}$ approximation algorithm unless $P = NP$.*

Proof. Such an algorithm could be used to approximate Independent Set in bounded-degree graphs. Given a graph with maximum degree d its inductive independence number ρ is also at most d . Trevisan [29] shows that there is no $d/2^{\mathcal{O}(\sqrt{\log d})}$ -approximation algorithm for all $d = \mathcal{O}(\log n)$ unless $P = NP$. This directly yields the claim. \square

As a second result we can also prove the impact of the number of channels k has to be as large as \sqrt{k} .

Theorem 6. *Even for $\rho = 1$ there is no $k^{\frac{1}{2}-\varepsilon}$ -approximation algorithm unless $ZPP = NP$.*

Our framework extends general combinatorial auctions with k items, and this is a standard result in the area [28, Chapter 9] derived from the hardness of independent set in general graphs.

In conclusion, our algorithmic results are supported by almost matching lower bounds in each parameter. Without further restricting the graph properties (which means to use additional properties of an interference model) no vitally better approximation guarantees can be achieved in terms of ρ resp. k . However, this does not prove no $\mathcal{O}(\rho + \sqrt{k})$ approximation can exist.

3 Edge-weighted Conflict Graphs

In this section we extend conflicts over binary relations (conflict/no-conflict). In wireless communication, we encounter situations that a radio transmission is exposed to interference by a number of devices relatively far away. If there was only a single one of them, interference would be acceptable but their overall interference is too high. For such aggregation aspects we introduce edge-weighted conflict graphs, in which there is a non-negative weight $w(u, v)$ between any pair of vertices $u, v \in V$. An *independent set* is defined as a set $M \subseteq V$ such that $\sum_{u \in M} w(u, v) < 1$ for all $v \in V$.

The definition of the inductive independence number can be generalized in a straightforward way. Since edge weights need not be symmetric, it turns out to be convenient to use the following symmetric edge weights $\bar{w}(u, v) = w(u, v) + w(v, u)$.

Definition 2. *The inductive independence number of an edge-weighted graph G is the minimum number ρ such that there is a total ordering $\pi: V \rightarrow [n]$ (bijective function) which fulfills for all vertices v and all independent sets $M \subseteq \{u \in V \mid \pi(u) < \pi(v)\}$ the following condition:*

$$\sum_{u \in M} \bar{w}(u, v) \leq \rho .$$

In the same way as in the unweighted case, we can use the definition to formulate the LP relaxation.

$$\text{Max. } \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x_{v,T} \quad (4a)$$

$$\text{s. t. } \sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T \subseteq [k] \\ j \in T}} \bar{w}(u, v) \cdot x_{u,T} \leq \rho \quad \text{for all } v \in V, j \in [k] \quad (4b)$$

$$\sum_{T \subseteq [k]} x_{v,T} \leq 1 \quad \text{for all } v \in V \quad (4c)$$

$$x_{v,T} \geq 0 \quad \text{for all } v \in V, T \subseteq [k] \quad (4d)$$

In weighted graphs we lose an important property we made use of in unweighted graphs: Resolving conflicts in one direction only does not suffice. To cope with this issue, we increase the scaling by another factor of 2. We use rounding and conflict resolution as previously to ensure that for each vertex v the sum of edge weights to neighboring vertices that have smaller indices and share a channel with v is at most $1/2$. Formally, a *partly-feasible allocation* is an allocation $S: V \rightarrow 2^{[k]}$ such that

$$\sum_{\substack{u \in V \\ \pi(u) < \pi(v) \\ S(v) \cap S(u) \neq \emptyset}} \bar{w}(u, v) < \frac{1}{2} \quad \text{for all } v \in V. \quad (5)$$

Rounding LP solutions to such partly-feasible allocations can be carried out in a similar way as Algorithm 1. Algorithm 2 decomposes the given LP solution the same way as Algorithm 1. Afterwards, it also performs two stages. In the rounding stage (lines 2–4), again a tentative allocation is determined randomly by considering the LP solution as a probability distribution.

Afterwards, only a partial conflict resolution (lines 5–8) is performed: If for some vertex v the sum of edge weights to neighbors that have lower π values and share a channel exceeds $1/2$, it is removed from the solution (i. e. it is allocated the empty set). Such a partly-feasible solution satisfies Equation (5).

Lemma 7. *For any feasible LP solution x^* with value b^* , Algorithm 2 calculates a partly-feasible allocation S of value at least $b^*/16\sqrt{k}\rho$ in expectation.*

Proof. The allocation is partly feasible since both allocations $S^{(1)}$ and $S^{(2)}$ satisfy Condition (5).

For the value of the solution let us again bound the value of the partly-feasible allocations $S^{(1)}$ and $S^{(2)}$. Again, let us fix $l \in \{1, 2\}$. Let $X_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the rounding stage. This time, we have

$$\mathbf{E}[X_{v,T}] = \frac{x_{v,T}^{(l)}}{4\sqrt{k}\rho}. \quad (6)$$

Let $X'_{v,T}$ be a 0/1 random variable indicating if $S^{(l)}(v)$ is set to T after the partial conflict-resolution stage. Again, we consider the event that $X'_{v,T} = 0$, given that $X'_{v,T} = 1$, i. e., that v is removed in the conflict-resolution stage after having survived the rounding stage.

Algorithm 2: LP rounding algorithm for the combinatorial auction problem with weighted conflict graphs

```

1 decompose  $x$  into two solutions  $x^{(1)}$  and  $x^{(2)}$  by  $x_{v,T}^{(1)} = x_{v,T}$  if  $|T| \leq \sqrt{k}$  and  $x_{v,S}^{(1)} = 0$ 
   otherwise.  $x^{(2)} = x - x^{(1)}$ ;
2 for  $l \in \{1, 2\}$  do
3   for  $v \in V$  do /* Rounding Stage */
4      $\lfloor$  with probability  $\frac{x_{v,T}}{4\sqrt{k}\rho}$  set  $S^{(l)}(v) := T$ 
5   for  $v \in V$  do /* Partial Conflict-Resolution Stage */
6     set  $U(v) := \{u \in V \mid \pi(u) < \pi(v),$ 
7        $S^{(l)}(v) \cap S^{(l)}(u) \neq \emptyset\}$ ;
8     if  $\sum_{u \in U(v)} \bar{w}(u, v) \geq \frac{1}{2}$  then
9        $S^{(l)}(v) := \emptyset$ 
9 return the better one of the allocations  $S^{(1)}$  and  $S^{(2)}$ 

```

Lemma 4 cannot be directly applied in this case. However, we have

$$\begin{aligned}
\Pr [X'_{v,T} = 0 \mid X_{v,T} = 1] &\leq \Pr \left[\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} \bar{w}(u, v) \cdot X_{u,T'} \geq \frac{1}{2} \right] \\
&\leq 2\mathbf{E} \left[\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} \bar{w}(u, v) \cdot X_{u,T'} \right] \\
&= \mathbf{E} \left[\sum_{\substack{u \in V \\ \pi(u) < \pi(v)}} \sum_{\substack{T' \subseteq [k] \\ T \cap T' \neq \emptyset}} \bar{w}(u, v) \cdot 2X_{u,T'} \right]
\end{aligned}$$

due to the Markov inequality and linearity of expectation. Now, we can use exactly the same arguments as in the proof of Lemma 4 literally. This proof relies on two conditions: the bound on $\mathbf{E}[X_{u,T}]$ and the fact that $x^{(l)}$ is a feasible LP solution. Both conditions are again satisfied.

This implies that $\Pr [X'_{v,T} = 0 \mid X_{v,T} = 1] \leq 1/2$ for both cases $l \in \{1, 2\}$. In combination with Equation (6), we get for all $v \in V$, $T \subseteq [k]$

$$\Pr [X'_{v,T} = 1] \geq \frac{x_{v,T}^{(l)}}{8\sqrt{k}\rho}$$

Thus, we can conclude that for $l \in \{1, 2\}$, we have

$$\mathbf{E} [b(S^{(l)})] \geq \frac{1}{8\sqrt{k}\rho} \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} x_{v,T}^{(l)} .$$

The expected value of the output is at least

$$\mathbf{E} \left[\max\{b(S^{(1)}), b(S^{(2)})\} \right] \geq \frac{1}{16\sqrt{k}\rho} \sum_{v \in V} \sum_{T \subseteq [k]} b_{v,T} \left(x_{v,T}^{(1)} + x_{v,T}^{(2)} \right) = \frac{b^*}{16\sqrt{k}\rho} .$$

□

Given a partly-feasible allocation S , Algorithm 3 implements the necessary additional conflict resolution to derive a fully-feasible one. The algorithm decomposes the partly-feasible allocation to a number of feasible candidate allocations S_1, S_2, \dots . Each allocation S_i is initialized such that $S_i(v) = S(v)$ if vertex v has been removed from all previous allocations S_1, \dots, S_{i-1} . Otherwise $S_i(v) = \emptyset$. Then a conflict resolution is performed on S_i : The vertices are considered by decreasing indices in the π ordering. If the weight bound is violated for some vertex v in the current allocation S_i , it is removed from the allocation by allocating the empty set. At the end, the best one of the candidate allocations is returned. We will see that each candidate allocation allocates at least half of the remaining vertices a non-empty set. Therefore at most $\lceil \log n \rceil$ candidates are computed and the best one has value at least $b(S)/\lceil \log n \rceil$.

Algorithm 3: Making a partly-feasible allocation fully feasible

```

1  $i := 1$  ;
2  $V' := V$  ;
3 while  $V' \neq \emptyset$  do
4   initialize  $S_i$  by  $S_i(v) := S(v)$  for  $v \in V'$  and  $S_i(v) := \emptyset$  otherwise ;
5   for  $v \in V'$  in order of decreasing  $\pi$  values do
6     if  $\sum_{u \in V', S_i(v) \cap S_i(u) \neq \emptyset} \bar{w}(u, v) < 1$  then
7       delete  $v$  from  $V'$                                 /*  $v$  stays in  $S_i$  */;
8     else
9        $S_i(v) := \emptyset$                                 /*  $v$  is removed from  $S_i$  */;
10   $i := i + 1$ ;
11 return the best one of the allocations  $S_1, S_2, \dots$ 

```

Lemma 8. *Given a (not necessarily feasible) allocation S in which Condition 5 is fulfilled for all $v \in V$, Algorithm 3 calculates a feasible allocation of value at least $b(S)/\lceil \log n \rceil$.*

Proof. Obviously, by construction all candidates are feasible and so is the output allocation.

Next, we prove that we need at most $\log n$ iterations of the while loop by showing that in each iteration at most half of the remaining vertices are removed from the allocation. This means the cardinality of V' is at least halved in each iteration. Let V'_i be the set V' after the i th iteration of the while loop; $V'_0 = V$.

Let us fix $i \in \mathbb{N}$, and $v \in V'_{i+1}$. We know that v has been removed from S_i by the algorithm. This only happens if

$$\sum_{\substack{u \in V'_i \\ S'_i(v) \cap S'_i(u) \neq \emptyset}} \bar{w}(u, v) \geq 1$$

where S'_i is the current state of S_i while the algorithm considers v . Since Equation 5 is obviously also satisfied for S'_i , it has to be

$$\sum_{\substack{u \in V'_i \\ \pi(u) > \pi(v) \\ S'_i(v) \cap S'_i(u) \neq \emptyset}} \bar{w}(u, v) \geq \frac{1}{2} .$$

For a vertex $u \in V'_i$ with $\pi(u) > \pi(v)$ we know that the vertex has either been removed from the allocation before (then $u \in V'_{i+1}$) or it stays in S_i (i. e. $S'_i(u) = S_i(u) = S(u)$ and $u \notin V'_{i+1}$). Hence

$$S'_i(u) = \begin{cases} \emptyset & \text{if } u \in V'_{i+1} \\ S(u) & \text{else} \end{cases} .$$

Combining these two insights, we get a necessary condition: if $v \in V'_{i+1}$ then

$$\sum_{u \in U_i(v) \setminus U_{i+1}(v)} \bar{w}(u, v) \geq \frac{1}{2} ,$$

where $U_i(v) = \{u \in V'_i \mid \pi(u) > \pi(v), S(v) \cap S(u) \neq \emptyset\}$. Summing up all $v \in V'_{i+1}$ we get

$$\sum_{v \in V'_{i+1}} \sum_{u \in U_i(v) \setminus U_{i+1}(v)} \bar{w}(u, v) \geq \frac{1}{2} |V'_{i+1}| .$$

On the other hand, we can change the ordering of the sums and use the symmetry of the weights \bar{w} to get

$$\begin{aligned} \sum_{v \in V'_{i+1}} \sum_{u \in U_i(v) \setminus U_{i+1}(v)} \bar{w}(u, v) &= \sum_{u \in V'_i \setminus V'_{i+1}} \sum_{\substack{v \in V'_{i+1} \\ \pi(u) > \pi(v) \\ S(v) \cap S(u) \neq \emptyset}} \bar{w}(u, v) \\ &= \sum_{v \in V'_i \setminus V'_{i+1}} \sum_{\substack{u \in V'_{i+1} \\ \pi(u) < \pi(v) \\ S(v) \cap S(u) \neq \emptyset}} \bar{w}(u, v) \\ &< \frac{1}{2} |V'_i \setminus V'_{i+1}| \end{aligned}$$

where the last bound is due to Condition (5).

In combination this yields

$$|V'_{i+1}| < |V'_i \setminus V'_{i+1}| ,$$

which implies

$$|V'_{i+1}| < \frac{1}{2} |V'_i| ,$$

meaning less than half of the remaining vertices are removed in each iteration.

So, since $|V'_0| = n$, we can conclude that

$$|V'_i| < \frac{1}{2^i} \cdot n .$$

Therefore, we get $|V'_{\lceil \log n \rceil}| < 1$. Thus the algorithm terminates within $\lceil \log n \rceil$ steps.

By definition, for all vertices $S_i(v) = S(v)$ for exactly one $i \in [\lceil \log n \rceil]$ and $S_i(v) = \emptyset$ else. So $\sum_{i \in [\lceil \log n \rceil]} b(S_i) = b(S)$. This yields for the value of the output

$$\max_{i \in [\lceil \log n \rceil]} b(S_i) \geq \frac{1}{\lceil \log n \rceil} \sum_{i \in [\lceil \log n \rceil]} b(S_i) = \frac{b(S)}{\lceil \log n \rceil}.$$

□

As a consequence, the computed feasible allocation has a value that in expectation is at most an $\mathcal{O}(\sqrt{k\rho \log n})$ factor smaller than that of the optimal LP solution.

4 Applications

In the previous sections we have described a general algorithmic approach to channel allocation problems when the underlying conflict graph has bounded inductive independence number. Here we will show that this property is particularly wide-spread among models for interference in wireless communication. Our aim is not to prove optimal bounds in each case but to show why we believe a bounded inductive independence number to be a key insight for understanding algorithmic problems in wireless networking.

The concept of conflict graphs can be applied in two basic scenarios. On the one hand, the task could be to allocate channels to *transmitters*. Each transmitter intends to cover a certain area, e.g., a base station in a cellular network. The interference model defines which transmitters can be assigned the same channels. On the other hand, instead of single transmitters one can consider pairs of network nodes (*links*) that act as sender and receiver. In such a scenario, “users” are no single network nodes but links. Therefore, the vertices of the conflict graph are links, and edges define which links can be assigned the same channels.

4.1 Transmitter Scenarios

A very simple, yet instructive model for a transmitter scenario is as follows. We have n transmitters located in the plane at points $p_1, \dots, p_n \in \mathbb{R}^2$. Each of the transmitters has a transmission range $r_1, \dots, r_n \in \mathbb{R}_{>0}$. Transmitters may be assigned the same channel if their transmission ranges do not intersect. Under these conditions interference constraints can be modeled by a disk graph. There is an edge between two vertices if the transmission-range disks around the corresponding receivers intersect.

Proposition 9. *Disk graphs have an inductive independence number of $\rho \leq 5$.*

Proof. Let $G = (V, E)$ be a disk graph. Let π be the ordering of vertices by decreasing radius of the corresponding disk. So in other words $V = \{v_1, \dots, v_n\}$ with $r_1 \geq r_2 \geq \dots \geq r_n$, where r_i is the radius of the disk around v_i . If the disk representation is given, this ordering can be computed in polynomial time by simply sorting the vertices. Let be $v \in V$ and $M \subseteq \{u \in V \mid \pi(u) < \pi(v)\}$ be an independent set in G . By definition of the ordering, we have for all $u \in M$ the radius is at least $r_{\pi(v)}$.

In order to show $|M| \leq 5$, we assume $|M| \geq 6$. This would yield that there were two vertices whose angle seen from v was at most 60° . From simple geometric arguments we can conclude

there has to be an edge between these two vertices. This contradicts the assumption that M is an independent set and thereby proves the claim. \square

Another example for the transmitter scenario is the so-called *distance-2 coloring*. In contrast to the above model not only the neighbors (with intersecting disks) must be on different channels but also their neighbors. Distance-2 coloring is a common model of transmitter scenarios. Here, we analyze the restriction on two graph classes. We refer the reader to [26] for the exact definitions and a discussion of the model. We can prove $\rho = \mathcal{O}(1)$ as well in this case.

Lemma 10. *Let $r > 0$, $a > 0$ and D be a disk of radius ar . Then the number of disks of radius at least r that intersect D but not each other is at most $(a + 2)^2$.*

Proof. W.l.o.g., we assume the surrounding disks to have radius exact r . By scaling them down and moving them inside their original area, they still do not intersect each other. By moving them to the respective closest location to D , they still intersect D .

The disks of radius r are fully contained within the disk of radius $kr + 2r$ around the center of D . Each takes an area of πr^2 , whereas the available area is only $\pi(ar + 2r)^2$. So, the number of surrounding disks is at most $\pi(ar + 2r)^2 / \pi r^2 = (a + 2)^2$. \square

Proposition 11. *For Distance-2 coloring in disk graphs the associated conflict graph has an inductive independence number $\rho = \mathcal{O}(1)$.*

Proof. As for disk graphs, we order the vertices by decreasing ranges. Now consider a vertex v and a conflicting vertex u of larger range. This vertex can either be directly connected to v (there are at most 5 ones of this kind) or via an intermediate vertex u' . If the $r_{u'} < r_v$ is smaller, we see that the disk of radius r_u around u intersects the one of radius $2r_v$ around v . The above lemma yields that there can be at most a constant number of such vertices. For the case $r_{u'} \geq r_v$, we take into consideration the disks around the intermediate vertices also do not intersect. So, there can be at most 5 intermediate vertices and as many conflicting vertices. The total number of conflicting vertices is constant. \square

Proposition 12. *For Distance-2 coloring in (r, s) -civilized graphs the inductive independence number of the associated conflict graph is $\rho \leq (4r/s + 2)$.*

Proof. In this case, the ordering does not matter. Therefore, we do not need to know the geometric representation of the graph.

Consider a vertex v and a set of vertices M conflicting with v but not with each other. Since the path length from v to each vertex in M is at most 2, the distance in the plane is at most $2r$. Now consider disks around the vertices in M , each of radius $s/2$. By definition of the (r, s) -civilized graph these disks do not intersect each other. However, each of them intersects a disk of radius $2r$ around v . Applying the above lemma, we see there are at most $(4r/s + 2)^2$ such disks. \square

As a matter of fact ρ has to depend on this ratio of r and s . Obviously, all graphs can be represented as (r, s) -civilized if the ratio r and s is unbounded. However, our algorithm's running time does not depend on r and s . Therefore, the approximation factor has to depend on them.

4.2 Unweighted Link-Based Scenarios

There are a number of different interference models for link-based scenarios that can be described by some unweighted conflict graph. They are often called graph-based interference models, but to avoid ambiguities we refer to them as *binary interference models*. Due to the large variety, we have to confine ourselves to some selected examples.

Probably the best known binary model is the *Protocol Model* [17]. Network nodes are modeled by points located in the plane. A link consisting of sender s and receiver r may be allocated to a channel if and only if for all other senders s' on this channel $d(s', r) \geq (1 + \Delta)d(s, r)$ for some constant $\Delta > 0$.

Proposition 13 (Wan [31]). *For the protocol model, the resulting conflict graph has an inductive independence number of*

$$\rho \leq \left\lceil \pi / \arcsin \frac{\Delta}{2(\Delta + 1)} \right\rceil - 1 .$$

The *IEEE 802.11 Model* by Alicherry et al. [2] is a bidirectional variant of this model, and in this case $\rho \leq 23$ [31].

A more graph-theoretical approach is *distance-2 matching* [3]. In this case, two edges $e \neq e'$ may be allocated to the same channel if there are at least two edges on any connecting path. Typically, results are restricted to certain graph classes, because in general approximating maximum distance-2 matchings is hard. For disk graphs, we can also show that the corresponding conflict graph has $\rho = \mathcal{O}(1)$. Interestingly, for distance-2 matching there is already an algorithm and analysis based on the observation that the inductive independence number is bounded, but the concepts are termed differently. Barrett et al. [4] analyze a greedy approach to find a maximum independent set. For a link $e = (u, v)$, they define $r(e) = r(u) + r(v)$, where $r(u)$ and $r(v)$ are the radius of the disk surrounding u resp. v . The algorithm orders the links by increasing values of $r(e)$. The key observation is now that for all links e the maximum number of links of higher index that collide with e but not with each other is $\mathcal{O}(1)$. This immediately yields $\rho = \mathcal{O}(1)$.

Corollary 14. *For distance-2 matching in disk graphs the associated conflict graph has an inductive independence number $\rho = \mathcal{O}(1)$.*

Analyses of greedy algorithms are often carried out in a similar manner. Such arguments already suffice to bound the inductive independence number. There is plenty of opportunity to further extend our results by similar observations.

4.3 Physical Model

The models mentioned above go well with graph-theoretic concepts. However, radio transmissions typically decrease asymptotically with increasing distance. The *physical model* captures this property much more accurately and is particularly wide-spread among engineers. Even though the physical model does not fit in the traditional binary graph-theoretic context, it has similar properties allowing it to be expressed using edge-weighted conflict graphs.

In this model, network nodes are located in a metric space. The received signal strength decreases as the distance increases. If a node transmits at a power level p , the signal strength at a distance of d is p/d^α , for a constant α . A transmission is received successfully if ratio of the received signal strength of the intended transmission and the strengths of concurrent transmissions

plus ambient noise is above some constant threshold $\beta > 0$. More formally, given pairs of senders s_i and receivers r_i that transmit at power level p_i , receiver r_i can decode the signal from sender s_i successfully if the SINR constraint

$$\frac{p_i}{d(s_i, r_i)^\alpha} \geq \beta \left(\sum_{j \in M \setminus \{i\}} \frac{p_j}{d(s_j, r_i)^\alpha} + \nu \right)$$

is fulfilled. Here M is the set of links transmitting at the same time on the same channel and $\nu \geq 0$ is a constant expressing ambient noise.

Note that we can easily reduce the model to a conflict graph if transmission powers are fixed. Prominent and simple classes of power assignments $p: V \rightarrow \mathbb{R}_{>0}$ are uniform ($p(v) = 1$) or linear ($p(v) = d(s_v, r_v)^\alpha$) assignments. More generally, we can consider assignments satisfying the following monotonicity constraints. If $d(\ell) \leq d(\ell')$ for two links ℓ, ℓ' then

$$p(\ell) \leq p(\ell') \quad \text{and} \quad \frac{p(\ell)}{d(\ell)^\alpha} \geq \frac{p(\ell')}{d(\ell')^\alpha} .$$

We furthermore assume the noise to play a minor role (cf. [25]).

Proposition 15. *The interference constraints in the physical model with fixed transmission power can be represented by a weighted conflict graph. If the power assignment satisfies the above constraints, the resulting inductive independence number is at most $\mathcal{O}(\log n)$.*

Proof. We choose the edges of the conflict graph to have the following weights. For $\ell = (s, r)$, $\ell' = (s', r')$ we set

$$w(\ell', \ell) = \min \left\{ 1, \frac{\beta}{1 + \varepsilon} \cdot \frac{p(\ell')}{d(s', r)^\alpha} \middle/ \left(\frac{p(\ell)}{d(s, r)^\alpha} - \frac{\beta}{1 + \varepsilon} \nu \right) \right\} ,$$

where

$$\varepsilon = \frac{\beta}{2} \min_{\ell=(s,r)} \min_{\ell'=(s',r')} \frac{p(\ell)}{d(s', r)^\alpha} \middle/ \frac{p(\ell)}{d(s, r)^\alpha} .$$

By this definition a set M fulfills the SINR constraint iff it corresponds to an independent set in the edge-weighted graph. The $1/(1+\varepsilon)$ factor is only necessary to get an exact transformation of the “ \geq ” in the SINR condition to the “ $<$ ” in the independent set definition. Apart from this factor the edge weights are equal to the notion of *affectance* a_p in [25], for which we have the following result.

Lemma 16 ([25]). *Let p be a power assignment satisfying Conditions 1 and 2 in [25]*

If M is a set of links that can concurrently transmit and ℓ is link with $d(\ell) \leq d(\ell')$ for all $\ell' \in M$, then

$$\sum_{\ell' \in M} a_p(\ell', \ell) = \mathcal{O}(1) \quad \text{and} \quad \sum_{\ell' \in M} a_p(\ell, \ell') = \mathcal{O}(\log n) .$$

This immediately yields the edge-weighted graph to have an inductive independence number $\rho = \mathcal{O}(\log n)$. \square

Interestingly, we can also use our approach if transmission powers are not given upfront. In this case, our algorithm has to decide about the assignment of links to channels and which transmission powers to use for each link. The first part is solved by LP rounding as above. In the LP we use edge weights ensuring that there is a feasible power assignment for the computed set of links. The second task of power assignment can then be done using a power control procedure by Kesselheim [24].

Note that, in contrast to the interference models mentioned above, in this case not all feasible solutions (i.e., feasibly scheduled sets of links) correspond to independent sets in the weighted graph. However, for our argument it suffices to observe that each set of feasible links corresponds to an LP solution for some ρ and that integral LP solutions with $\rho = 1$ also correspond to feasible sets of links. Combining these insights with the bounds in [24] and the ones we proved above, we obtain the following result.

Theorem 17. *There is a choice of edge weights such that our algorithm in combination with the power control procedure in [24] achieves an $\mathcal{O}(\sqrt{k} \log n)$ approximation in fading metrics and an $\mathcal{O}(\sqrt{k} \log^2 n)$ approximation in general metrics.*

Proof. We define the weighted graph as follows. The set of vertices is again the set of all links \mathcal{R} . The ordering π is the ordering from large to small distances between the sender and its receiver. Between two links $\ell = (s, r)$ and $\ell' = (s', r')$, we have the following weight

$$w(\ell, \ell') = \begin{cases} \frac{1}{\tau} \min \left\{ 1, \frac{d(s, r)^\alpha}{d(s, r')^\alpha} \right\} + \frac{1}{\tau} \min \left\{ 1, \frac{d(s, r)^\alpha}{d(s', r)^\alpha} \right\} & \text{if } \pi(\ell) < \pi(\ell') \\ 0 & \text{otherwise} \end{cases},$$

$$\text{where } \tau = \frac{1}{2 \cdot 3^\alpha \cdot (4\beta + 2)}.$$

Theorem 3 in [24] states that for each independent set in the weighted graph the power control algorithm calculates a feasible set of links.

On the other hand Theorem 1 in [24] shows that under the above edge weights each feasible set of links is also an LP solution for some $\rho = \mathcal{O}(1)$ in fading metrics. Theorem 7 in [24] shows $\rho = \mathcal{O}(\log n)$ in general metrics.

In conclusion, this implies that by applying our rounding algorithm to the LP using above defined weights we get a solution, for which we can apply the power assignment of [24] to obtain a feasible set of links. The resulting allocation is an $\mathcal{O}(\sqrt{k} \log n)$ approximation for fading metrics and an $\mathcal{O}(\sqrt{k} \log^2 n)$ approximation in general metrics. \square

5 Mechanism Design

In this section we show how to apply the framework proposed by Lavi and Swamy [27] to obtain a truthful mechanism for the problem, in which the valuations for the allocations are private information. We only highlight the main ideas of this technique and the most important observations that allow the use for our problem.

The main idea of the approach is to decompose an optimal LP solution x^* into a set of polynomially many integral solutions with the following property. For each integral solution we determine a probability, and the expected social welfare of a randomly chosen solution according to the probabilities is exactly b^*/α , where in our case $\alpha = 8 \cdot \sqrt{k} \cdot \rho$. Given such a decomposition, we can use

scaled VCG payments to implement a randomized mechanism that is truthful in expectation. For an accessible presentation of the complete technique, see [28, Chapter 12] or [27].

In particular, for simplicity let us first consider only a constant number of channels; the adjustment to arbitrary many channels is treated below. We ask the vertices to obtain all valuations for all channel bundles and solve the corresponding LP (interference information is assumed to be publicly available). Note that at this point we are given the optimal solution to an *infeasible* LP. We set up a decomposition LP with exponentially many variables – one for each *feasible* integral solution – that represent our desired probabilities. This LP has polynomially many constraints but exponentially many variables. We can construct the dual with polynomially many variables and exponentially many constraints. The variables can be interpreted as valuations in an adjusted combinatorial auction problem. If this problem has an algorithm that verifies an integrality gap, we obtain a separation oracle and can solve the dual decomposition LP in polynomial time. In particular, it allows us to construct an equivalent LP with a polynomial number of constraints, i.e., the ones corresponding to the solutions obtained by our algorithm. For this polynomial-sized dual we construct the primal and determine the polynomially many probabilities of the solutions found by our algorithm, which completes the decomposition.

It remains to verify that our algorithms provide integral solutions within the desired integrality gap of α for the adjusted combinatorial auction problems using dual variables as valuations. We note here that our algorithms bound the integrality gap of LP (1) and (4), and they can be derandomized using the technique of pairwise independence. In this way, given an optimal LP solution x^* we can obtain an integral solution of value at least b^*/α . Note that our LP describe, in fact, relaxations of the combinatorial auction problem with conflict graphs, because Conditions (1b) and (4b) allow each vertex to have multiple neighbors on the same channel. An arbitrary integral solution to the LP might thus be infeasible for the original problem. This is even more severe in the case of the physical model with power control, where even the interpretation of edge weights is significantly disconnected from the actual interference that is received. However, our algorithms produce feasible integral solutions with the desired gap to the infeasible fractional optimum. Thus, they also prove the gap for a potential fractional optimum to the LP describing the (more constrained) exact combinatorial auction problem with conflict graphs in the respective cases. The remaining arguments can be adapted from [27] almost without adjustment.

In case of an arbitrary number of channels, we can use demand oracles to solve the LPs. This results in only a polynomial number of (non-zero) variables for the LP and for the dual of the decomposition LP. Note that the procedure to separate the dual of the decomposition LP does not require demand oracles. In fact, the complete decomposition procedure can be carried out without accessing the original bidder valuations.

6 Asymmetric Channels

Up to now, channels were symmetric in terms of interference, which means the same interference model is applied to each channel. In a more general setting, for each of the k channels a different edge set E_j resp. a different edge-weight function w_j for the interference graph is given.

In this case, we have an edge weight function \bar{w}_j for each channel $j \in [k]$. The above LP relaxation be easily adapted by exchanging \bar{w} by \bar{w}_j in the constraints (1b). In contrast, the analysis of the rounding algorithms internally depends on the assumption of symmetric channels. In particular, the proof of Lemma 4 uses the symmetry.

However, when exchanging the probability for a vertex v to choose set T by $x_{v,T}^{(l)}/2k\rho$ resp. $x_{v,T}^{(l)}/4k\rho$, the proof of Lemma 4 can be carried out the same way without using the symmetry.

Hence, for the asymmetric case, we lose a factor of $\mathcal{O}(k \cdot \rho)$ resp. $\mathcal{O}(k \cdot \rho \cdot \log n)$ in the LP rounding step. This represents our approximation ratio. The result may seem like a trivial generalization of the $k = 1$ case. However, this is not true as multiple graphs make the problem much harder. We can justify the approximation factor by a hardness bound.

Theorem 18. *For each ρ, k with $\rho \cdot k = \mathcal{O}(\log n)$ there is no $\rho \cdot k / 2^{\mathcal{O}(\sqrt{\log(\rho \cdot k)})}$ approximation algorithm for asymmetric channels unless $P = NP$.*

Proof. Again, such an algorithm could be used to approximate the independent set problem in bounded-degree graphs. Given a graph $G = (V, E)$ with maximum degree d , we construct k graphs $G_1 = (V, E_1), \dots, G_k = (V, E_k)$ each having an inductive independence number of $\rho = d/k$. For simplicity of notation, we assume this is an integer.

Let $\{v_1, \dots, v_n\}$ be an arbitrary ordering of the vertices. We now distribute the edges from E to the edge sets E_1, \dots, E_k . For a vertex v_i the incident edges to vertices v_j of lower index are distributed such that each edge set gets at most ρ such edges. Since the maximum vertex degree is d this is always possible. The valuations for the vertices are chosen such that for all vertices $b_{v,T}$ is 1 only for $T = [k]$ and 0 otherwise.

By this construction allocations of valuation b exactly correspond to independent sets of size b . Thus, such an approximation algorithm cannot exist unless $P = NP$. \square

As we see, for asymmetric channels our algorithms are close to optimal without making further assumptions about the interference model.

7 Open Problems

In this paper we present a general framework for secondary spectrum auctions that works with a large number of interference models. Our approach can easily be extended to even more models by proving bounds on the inductive independence number in the associated graphs. To improve the results in this paper, it would, e.g., be interesting to know if for the physical model it also holds that $\rho = \mathcal{O}(1)$ in general metrics or for distance-based power assignments.

For obtaining a truthful mechanism we use decomposition and rounding of LP solutions, and we heavily rely on the ellipsoid method. It is an interesting question if this could be avoided to make the algorithm more applicable in practice.

References

- [1] Karhan Akcoglu, James Aspnes, Bhaskar DasGupta, and Ming-Yang Kao. Opportunity cost algorithms for combinatorial auctions. *CoRR*, cs.CE/0010031, 2000.
- [2] Mansoor Alicherry, Randeep Bhatia, and Li (Erran) Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *Proceedings of the 11th International Conference on Mobile Computing and Networking (MobiCom)*, pages 58–72, 2005.

- [3] H. Balakrishnan, C.L. Barrett, V.S.A. Kumar, M.V. Marathe, and S. Thite. The distance-2 matching problem and its relationship to the mac-layer capacity of ad hoc wireless networks. *Selected Areas in Communications, IEEE Journal on*, 22(6):1069 – 1079, aug. 2004.
- [4] Christopher L. Barrett, V. S. Anil Kumar, Madhav V. Marathe, Shripad Thite, and Gabriel Istrate. Strong edge coloring for channel assignment in wireless radio networks. In *Proceedings of the 4th annual IEEE international conference on Pervasive Computing and Communications Workshops (PERCOMW)*, page 106, 2006.
- [5] Milind M. Buddhikot, Paul Kolodzy, Scott Miller, Kevin Ryan, and Jason Evans. Dimsum-net: New directions in wireless networking using coordinated dynamic spectrum access. In *Proceedings of the IEEE WoWMoM*, pages 78–85, 2005.
- [6] George Christodoulou, Khaled Elbassioni, and Mahmoud Fouz. Truthful mechanisms for exhibitions. In *Proceedings of the 6th Workshop on Internet & Network Economics (WINE)*, pages 170–181, 2010.
- [7] Peter Cramton, Yoav Shoham, and Richard Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.
- [8] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the 38th ACM Symposium on Theory of Computing (STOC)*, pages 644–652, 2006.
- [9] Alexander Fanghänel, Sascha Geulen, Martin Hoefer, and Berthold Vöcking. Online capacity maximization in wireless networks. In *Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 92–99, 2010.
- [10] Alexander Fanghänel, Thomas Kesselheim, Harald Räcke, and Berthold Vöcking. Oblivious interference scheduling. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 220–229, 2009.
- [11] Alexander Fanghänel, Thomas Kesselheim, and Berthold Vöcking. Improved algorithms for latency minimization in wireless networks. In *Proceedings of the 36th International EATCS Colloquium on Automata, Languages and Programming (ICALP)*, volume 2, pages 208–219, 2009.
- [12] Uriel Feige and Jan Vondrák. Approximation algorithms for allocation problems: Improving the factor of $1-1/e$. In *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 667–676, 2006.
- [13] Aleksei V. Fishkin. Disk graphs: A short survey. In *Proceedings of the First Workshop on Approximation and Online Algorithms (WAOA)*, pages 260–264, 2003.
- [14] S. Gandhi, C. Buragohain, Lili Cao, Haitao Zheng, and S. Suri. A general framework for wireless spectrum auctions. In *Proceedings of the 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 22–33, April 2007.
- [15] Olga Goussevskaya, Roger Wattenhofer, Magnús M. Halldórsson, and Emo Welzl. Capacity of arbitrary wireless networks. In *Proceedings of the 28th Conference of the IEEE Communications Society (INFOCOM)*, pages 1872–1880, 2009.

- [16] Albert Gräf, Martin Stumpf, and Gerhard Weißenfels. On coloring unit disk graphs. *Algorithmica*, 20:277–293, 1994.
- [17] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46:388–404, 2000.
- [18] Magnús M. Halldórsson. Wireless scheduling with power control. In *Proceedings of the 17th European Symposium on Algorithms (ESA)*, pages 361–372, 2009.
- [19] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999.
- [20] Dorit Hochbaum. Efficient bounds for the stable set, vertex cover and set packing problems. *Discrete Appl. Math.*, 6(3):243–254, 1983.
- [21] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004.
- [22] Omer Ileri, Dragan Samardzija, Theodore Sizer, and Narayan B. Mandayam. Demand responsive pricing and competitive spectrum allocation via a spectrum server. In *Proceedings of the 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 194–202, 2005.
- [23] Akihisa Kako, Takao Ono, Tomio Hirata, and Magnús M. Halldórsson. Approximation algorithms for the weighted independent set problem. In *WG*, pages 341–350, 2005.
- [24] Thomas Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1549–1559, 2011.
- [25] Thomas Kesselheim and Berthold Vöcking. Distributed contention resolution in wireless networks. In *Proceedings of the 24th Symposium on Distributed Computing (DISC)*, pages 149–163, 2010.
- [26] S.O. Krumke, M.V. Marathe, and SS Ravi. Models and approximation algorithms for channel assignment in radio networks. *Wireless Networks*, 7(6):575–584, 2001.
- [27] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 595–604, 2005.
- [28] Noam Nisan, Éva Tardos, Tim Roughgarden, and Vijay Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [29] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC)*, pages 453–461, 2001.
- [30] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th ACM Symposium on Theory of Computing (STOC)*, pages 67–74, 2008.

- [31] Peng-Jun Wan. Multiflows in multihop wireless networks. In *Proceedings of the 10th ACM International Symposium Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, pages 85–94, 2009.
- [32] Yuli Ye and Allan Borodin. Elimination graphs. In *Proceedings of the 36th International EATCS Colloquium on Automata, Languages and Programming (ICALP)*, pages 774–785, 2009.
- [33] Xia Zhou, Sorabh Gandhi, Subhash Suri, and Haitao Zheng. eBay in the Sky: Strategy-proof wireless spectrum auctions. In *Proceedings of the 14th International Conference on Mobile Computing and Networking (MobiCom)*, pages 2–13, 2008.
- [34] Xia Zhou and Haitao Zheng. TRUST: A general framework for truthful double spectrum auctions. In *Proceedings of the 28th Conference of the IEEE Communications Society (INFOCOM)*, pages 999–1007, 2009.